


#10
RWB
10-23-01

ADVERTISING EMPLOYING WATERMARKING

Related Application Data

5 This application is a continuation-in-part of copending application 09/130,624, filed August 6, 1998, which is a continuation of application 08/508,083 (now Patent 5,841,978).

This application is also a continuation-in-part of copending application 09/314,648, filed May 19, 1999 (attached as Appendix A). This application is also a continuation-in-part of copending provisional application 60/134,782, also filed May 19, 1999 (attached as Appendix B).

10

Field of the Invention

The present invention relates optical user interfaces that sense digitally-encoded objects. The invention further relates to systems using such optical interfaces to control computers, and to navigate over or act as portals on networks.

15

Background and Summary of the Invention

20 "Bedoop." That might be the sound that someone might hear as they lazily place a magazine advertisement in front of their desktop camera. Magically, the marketing and sales web site associated with the ad is displayed on their computer. More information? Want to buy now? Look at the full product line? No problem.

25 "Bedoop." That might be the same sound when that same someone places their credit card in front of their desktop camera. Instantly, the product displayed on the web page is purchased. Behind the scenes, a secure purchase link is initiated, transmitting all requisite information to the vendor. Twist the credit card clockwise and the purchaser chooses overnight delivery.

30 So goes an exemplary embodiment of the invention further described in this application. Though this example is rather specific, it nevertheless alludes to an indescribably vast array of applications possible when a digital camera or other optical sensing device is turned into a general purpose user interface device with an intuitive power that very well might rival the mouse and the keyboard.

The centerpiece of the invention is that an object or paper product so-scanned contains digital information that can be quickly read and acted upon by an appropriately configured device, computer or appliance. The preferred embodiment envisions that this digital information is aesthetically hidden on objects. These objects have been previously and pro-actively marked with the digital information, using any of the broad ranges of printing and processing techniques which are available on the market and which are widely described in the open literature and patent literature surrounding digital watermarking.

Be this as it may, though the invention concentrates on flat object applications wherein the digital information is often imperceptibly integrated into the object, it is certainly not meant to be so limited. Objects can be three dimensional in nature and the information more visually overt and/or pre-existing (i.e., not “pro-actively” embedded, or not even be “digital,” per se). Different implementation considerations attach to these variants. Likewise, though the bulk of this disclosure concentrates on objects which have some form of digital message attached thereto, some aspects of the invention may apply to objects which have no such thing, where the prior arts of pattern recognition and gestural input can be borrowed in combination with this invention to effect yet a broader array of applications.

“Bedoop.” The sound that a refrigerator might make, outfitted with a simple camera/processor unit/net connection, as the ten year old holds up the empty milk carton and a ping goes out to the local grocery store, adding the item to an accumulating delivery list. The sound that might be heard echoing over and over inside Internet cafés as heretofore computerphobes take their first skeptical steps onto the world wide web. The sound heard at the fast food counter as the repeat customer holds up their sandwich card ticking off their latest meal, hoping for the sirens to go off for a \$500 prize given to the lucky customer of the week. Blue sky scenarios abound.

This invention is therefore about powerful new user interfaces to computers involving optical input. These new user interfaces extend into the everyday world in ways that a mouse and keyboard never could. By enabling everyday objects to communicate their identities and functions to ever-attendant devices, not only will the world wide web be given an entirely new dimension, but basic home and office computing may be in store for some fundamental advances as well.

These and a great many other features of the present invention will be more readily apparent from the following detailed description, which proceeds with reference to the accompanying drawings.

Brief Description of the Drawings

Fig. 1 is a block diagram showing one embodiment of the present invention.

Fig. 2 is another block diagram showing an embodiment of the present invention.

Detailed Description

Basically, the technology detailed in this disclosure may be regarded as enhanced systems by which users can interact with computer-based devices. Their simple nature, and adaptability for use with everyday objects (e.g., milk cartons), makes the disclosed technology well suited for countless applications.

Due to the great range and variety of subject matter detailed in this disclosure, an orderly presentation is difficult to achieve. As will be evident, many of the topical sections presented below are both founded on, and foundational to, other sections. For want of a better rationale, the sections are presented below in a more or less random order. It should be recognized that both the general principles and the particular details from each section find application in other sections as well. To prevent the length of this disclosure from ballooning out of control, the various permutations and combinations of the features of the different sections are not exhaustively detailed. The inventors intend to explicitly teach such combinations/permutations, but practicality requires that the detailed synthesis be left to those who ultimately implement systems in accordance with such teachings.

Basic Principles – Refrigerators and Clutter

Referring to Fig. 1, a basic embodiment 10 of the present invention includes an optical sensor 12, a computer 14, and a network connection 16 to the internet 18. The illustrated optical sensor 12 is a digital camera having a resolution of 320 by 200 pixels (color or black and white) that stares out, grabbing frames of image data five times per second and storing same in one or more frame buffers. These frames of image data are analyzed by a computer 14 for the presence of Bedoop data. (Essentially, Bedoop data is any form of digital data

encoding recognized by the system 10 -- data which, in many embodiments, initiates some action.) Once detected, the system responds in accordance with the detected Bedoop data (e.g., by initiating some local action, or by communication with a remote computer, such as over the internet, via an online service such as AOL, or using point-to-point dial-up communications, as with a bulletin board system.

Consider the milk carton example. The artwork on a milk carton can be adapted to convey Bedoop data. In the preferred embodiment, the Bedoop data is steganographically encoded (e.g., digitally watermarked) on the carton. Numerous digital watermarking techniques are known -- all of which convey data in a hidden form (i.e., on human inspection, it is not apparent that digitally encoded data is present). Exemplary techniques operate by slightly changing the luminance, or contours, of selected points on artwork or text printed on the carton, or splatter tiny droplets of ink on the carton in a seemingly random pattern. Each of these techniques has the effect of changing the local luminance at areas across the carton -- luminance changes that can be detected by the computer 14 and decoded to extract the encoded digital data. In the case of a milk carton, the data may serve to identify the object as, e.g., a half gallon carton of Alpenrose brand skim milk.

The Fig. 1 apparatus can be integrated into the door of a refrigerator and used to compile a shopping list. Milk cartons, and other Bedoop-encoded packaging 20, can be held up the optical sensor. When the computer 14 detects the presence of Bedoop data and successfully decodes same, it issues a confirmation tone ("be-doop") from a speaker or other audio transducer 22. The computer then adds data identifying the just-detected object to a grocery list. This list can be maintained locally (in disk storage, non-volatile RAM 24, or the like in the refrigerator, or elsewhere in the home), or remotely (e.g., at a server located at a user-selected grocery, or elsewhere). In either event, the list is desirably displayed on a display in the user's home (e.g., an LCD screen 26 built into the front of the appliance). Conventional user interface techniques can be employed permitting the user to scroll through the displayed list, delete items as desired, etc.

Periodically, the listed groceries can be purchased and the list cleared. In one embodiment, the list is printed (either at the home or at the grocery), and the user walks the grocery aisles and purchases same in the conventional manner. In another embodiment, the grocer pulls the listed items from the shelves (in response to a user request conveyed by the

internet or telephone, or by a gesture as hereafter detailed). Once the list has been pulled, the grocer can alert the user that the groceries are available for pickup (again, e.g., by internet or telephone message), or the grocer can simply deliver the groceries directly to the user's home. Naturally, on-line payment mechanisms can be employed if desired.

5 Consider a wholly unrelated Bedoop application. An Excel spreadsheet is printed onto paper, and the paper becomes buried in a stack of clutter on an office worker's desk. Months later the spreadsheet again becomes relevant and is dug out of the stack. Changes need to be made to the data, but the file name has long-since been forgotten. The worker simply holds the dug-out page in front of a camera associated with the desktop computer. A moment later,
10 the electronic version of the file appears on the worker's computer display.

 When the page was originally printed, tiny droplets of ink or toner were distributed across the paper in a pattern so light as to be essentially un-noticeable, but which steganographically encoded the page with a plural-bit binary number (e.g., 64 bits). A database (e.g., maintained by the operating system, the Excel program, the printer driver, etc.)
15 stored part of this number (e.g., 24 bits, termed a Universal Identifier or UID) in association with the path and file name at which the electronic version of the file was stored, the page number within the document, and other useful information (e.g., author of the file, creation date, etc.).

 The steganographic encoding of the document, and the updating of the database, can
20 be performed by the software application (e.g., Excel). This option can be selected once by the user and applied thereafter to all printed documents (e.g., by a user selection on an "Options" drop-down menu), or can be presented to the user as part of the Print dialog window and selected (or not) for each print job.

 When such a printed page is later presented to the camera, the computer automatically
25 detects the presence of the encoded data on the page, decodes same, consults the database to identify the file name/location/page corresponding to the UID data, and opens the identified file to the correct page (e.g., after launching Excel). This application is one of many "paper as portal" applications of the Bedoop technology.

 The foregoing are but two of myriad applications of the technology detailed herein. In
30 the following discussion a great many other applications are disclosed (some groundbreaking,

a few gimmicky). However, regardless of the length of the specification, it is possible only to begin to explore a few of the vast ramifications of this technology.

A few more details on the basic embodiments described above may be helpful before delving into other applications.

5

Optics

For any system to decode steganographically encoded data from an object, the image of the object must be adequately focused on the digital camera's CCD (or other) sensor. In a low cost embodiment, the camera has a fixed nominal focal length, e.g., in the range of 6-24 inches (greater or lesser lengths can of course be used). Since the camera is continuously grabbing and analyzing frames of data, the user can move the object towards- or away- from the sensor until the decoder succeeds in decoding the steganographically encoded data and issues a confirming "Bedoop" audio signal.

10

In more elaborate embodiments, known auto-focusing technology can be employed.

15

In still other embodiments, the camera (or other sensor) can be equipped with one or more auxiliary fixed-focus lenses that can be selectively used, depending on the particular application. Some such embodiments have a first fixed focused lens that always overlies the sensor, with which one or more auxiliary lenses can be optically cascaded (e.g., by hinge or slide arrangements). Such arrangements are desirable, e.g., when a camera is not a dedicated Bedoop sensor but also performs other imaging tasks. When the camera is to be used for Bedoop, the auxiliary lens is positioned (e.g., flipped into) place, changing the focal length of the first lens (which may be unsuitably long for Bedoop purposes, such as infinity) to an appropriate Bedoop imaging range (such as one foot).

20

Other lens-switching embodiments do not employ a fixed lens that always overlies the sensor, but instead employ two or more lenses that can be moved into place over the sensor. By selecting different lenses, focal lengths such as infinity, six feet, and one foot can be selected.

25

In all such arrangements, it is desirable (but not essential) that the steganographically-encoded portion of the object being imaged fills a substantial part of the image frame. The object can be of various sizes, e.g., an 10 by 12 inch front panel of a cereal box, or a proof of purchase certificate that is just one inch square. To meet this requirement, small objects will

30

obviously need to be placed closer to the camera than large objects. The optics of the system can be designed, e.g., by selection of suitable aperture sizing and auxiliary lighting (if needed), to properly image objects of various sizes within a range of focal distances.

Some embodiments avoid issues of focal distances and identifying the intended object by constraining the size of the object and/or its placement. An example is a business card reader that is designed for the sole task of imaging business cards. Various such devices are known.

Decoding/Encoding

The analysis of the image data can be accomplished in various known ways. Presently, most steganographic decoding relies on general purpose microprocessors that are programmed by suitable software instructions to perform the necessary analysis. Other arrangements, such as using dedicated hardware, reprogrammable gate arrays, or other techniques, can of course be used.

The steganographic decoding process may entail three steps. In the first, the object is located. In the second, the object's orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object.

The first step, object location, can be assisted by various clues. One is the placement of the object; typically the center of the image field will be a point on the object. The surrounding data can then be analyzed to try and discern the object's boundaries.

Another location technique is slight movement. Although the user will typically try to hold the object still, there will usually be some jitter of the Bedoop object within the image frame (e.g., a few pixels back and forth). Background visual clutter, in contrast, will typically be stationary. Such movement may thus be sensed and used to identify the Bedoop object from within the image data.

Still another object-location clue is object shape. Many Bedoop objects are rectangular in shape (or trapezoidal as viewed by the camera). Straight edge boundaries can thus be used to define an area of likely Bedoop data.

Color is a further object identification clue that may be useful in some contexts.

Yet another object location clue is spatial frequency. In imaging systems with well defined focal zones, undesired visual clutter may be at focal distances that results in blurring.

The Bedoop object, in contrast, will be in focus and may be characterized by fine detail. Analyzing the image data for the high frequencies associated with fine detail can be used to distinguish the intended object from others.

Characteristic markings on the object (as discussed below in connection with determining object orientation), can also be sensed and used in locating the object.

Once the Bedoop object has been located within the image data, masking can be applied (if desired) to eliminate image data not corresponding to the intended object.

The second step in the decoding process - determining orientation of the Bedoop data - can likewise be discerned by reference to visual clues. For example, some objects include subliminal graticule data, or other calibration data, steganographically encoded with the Bedoop data to aid in determining orientation. Others can employ overt markings, either placed for that sole purpose (e.g. reference lines or fiducials), or serving another purpose as well (e.g. lines of text), to discern orientation. Edge-detection algorithms can also be employed to deduce the orientation of the object by reference to its edges.

Some embodiments filter the image data at some point in the process to aid in ultimate Bedoop data extraction. One use of such filtering is to mitigate image data artifacts due to the particular optical sensor. For example, CCD arrays have regularly-spaced sensors that sample the optical image at uniformly spaced discrete points. This discrete sampling effects a transformation of the image data, leading to certain image artifacts. An appropriately configured filter can mitigate the effect of these artifacts.

(In some arrangements, the step of determining the orientation can be omitted. Business card readers, for example, produce data that is reliably free of artifacts and is of known scale. Or the encoding of the Bedoop data can be effected in such a way that renders it relatively immune to certain distortion mechanisms. For example, while the presently-preferred encoding arrangement operates on a 2D grid basis, with rows and columns of data points, the encoding can alternatively be done on another basis (e.g., a rotationally-symmetric form of encoding, such as a 2D bar-code, so that rotational state of the image data can be ignored). In still other embodiments, the orientation-determining step can be omitted because the decoding can readily proceed without this information. For example decoding which relies on the Fourier-Mellin transform produces data in which scale and rotation can be ignored.)

Once the orientation of the object is discerned, the image data may be virtually re-registered, effectively mapping it to another perspective (e.g., onto a rectilinear image plane). This mapping can employ known image processing techniques to compensate, e.g., for rotation state, scale state, differential scale state, and X-Y offset, of the original Bedoop image data. The resulting frame of data may then be more readily processed to extract the steganographically-encoded Bedoop data.

In the preferred embodiment, after the image data is remapped into rectilinear planar form, subliminal graticule data is sensed that identifies the locations within the image data where the binary data is encoded. Desirably, the binary data is redundantly encoded, e.g., in 8 x 8 patch blocks. Each patch comprises one or more pixels. (The patches are typically square, and thus contain 1, 4, 9, or 16, etc. pixels.) The nominal luminance of each patch before encoding (e.g., artwork pre-existing on the object) is slightly increased or decreased to encode a binary "1" or "0." The change is slight enough to be generally imperceptible to human observers, yet statistically detectable from the image data – especially if several such blocks are available for analysis. Preferably, the degree of change is adapted to the character of the underlying image, with relatively greater changes being made in regions where the human eye is less likely to notice them. Each block thus encoded can convey 64 bits of data. The encoding of such blocks in tiled fashion across the object permits the data to be conveyed in robust fashion.

Much of the time, of course, the Bedoop sensor is staring out and grabbing image frames that have no Bedoop data. Desirably, the detection process includes one or more checks to assure that Bedoop data is not wrongly discerned from non-Bedoop image data. Various techniques can be employed to validate the decoded data, e.g., error detecting codes can be included in the Bedoop payload and checked to confirm correspondence with the other Bedoop payload. Likewise, the system can confirm that the same Bedoop data is present in different tiled excerpts within the image data, etc.

Pub 27 (Details of the preferred encoding techniques are further detailed in co-pending applications 09/293,601, filed April 15, 1999, entitled METHODS AND DEVICES FOR RECOGNIZING BANKNOTES AND RESPONDING ACCORDINGLY, 09/127,502, filed July 31, 1998, and patent 5,862,260.)

Data Structures, Formats, Protocols, and Infrastructures

In an exemplary system, the Bedoop data payload is 64 bits. This payload is divided into three fields CLASS (12 bits), DNS (24 bits) and UID (24 bits). (Other payload lengths, fields, and divisions, are of course possible, as is the provision of error-checking or error-correcting bits.)

Within the above-described eight patch - by - eight patch data block, the bits are ordered row by row, starting with the upper left patch. The first 12 bits are the CLASS ID, followed by 24 bits of DNS data followed by 24 bits of UID data. (In other embodiments, the placement of bits comprising these three fields can be scrambled throughout the block.)

Briefly, the CLASS ID is the most rudimentary division of Bedoop data, and may be analogized, in the familiar internet taxonomy, to the limited number of top level domains (e.g., .com, .net, .org, .mil, .edu, .jp, .de, .uk, etc.). It is basically an indicator of object type. The DNS ID is an intermediate level of data, and may be analogized to internet server addresses (e.g., biz.yahoo, interactive.wsj, etc.) The UID is the finest level of granularity, and can roughly be analogized to internet pages on a particular server (e.g., edition/current/summaries/front.htm, daily/home/default.hts, etc.).

Generally speaking, the CLASS ID and DNS ID, collectively, indicate to the system what sort of Bedoop data is on the object. In the case of Bedoop systems that rely on remote servers, the CLASS and DNS IDs are used in identifying the server computer that will respond to the Bedoop data. The UID determines precisely what response should be provided.

In the case of a refrigerator Bedoop system, what happens if an object with an unfamiliar CLASS/DNS ID data is encountered? The system can be programmed not to respond at all, or to respond with a raspberry-like sound (or other feedback) indicating "I see a Bedoop object but don't know what to do with it."

Most systems will be able to respond to several classes of Bedoop objects. Simple software-based systems can compare the CLASS/DNS ID (and optionally the UID) to fixed values, and can branch program execution to corresponding subroutines. Likewise, hardware-based systems can activate different circuitry depending on the detected CLASS/DNS ID.

In the case of a computer equipped with a Bedoop input device (e.g., a Sony VAIO PictureBook laptop with built-in camera), the operating system's registry database can be

employed to associate different application programs with different CLASS/DNS IDs (just as the .XLS and .DOC file extensions are commonly associated by existing operating system registries to invoke Microsoft Excel and Word software applications, respectively). When a new Bedoop application is installed, it logs an entry in the registry database indicating the CLASS/DNS ID(s) that it will handle. Thereafter, when an object with such a CLASS/DNS ID is encountered, the operating system automatically launches the corresponding application to service the Bedoop data in an appropriate manner.

Sometimes the computer system may encounter a Bedoop object for which it does not have a registered application program. In such case, a default Bedoop application can be invoked. This default application can, e.g., establish an internet link to a remote server computer (or a network of such computers), and can transmit the Bedoop data (or a part of the Bedoop data) to that remote computer. The remote server can undertake the response itself, it can instruct the originating computer how to respond appropriately, or it can undertake some combination of these two responses. (Such arrangements are further considered below.)

Fig. 2 shows an illustrative architecture employing the foregoing arrangement.

At a local Bedoop system 28 (which may be implemented, for example, using a conventional personal computer 29), a camera, scanner, or other optical sensor 30 provides image data to a decoder 32 (which may be implemented as a software component of the operating system 33). The decoder 32 analyzes the image data to discern the plural-bit Bedoop data. The CLASS ID of this Bedoop data is applied to a Bedoop registry 34. The registry responds by identifying and launching a local Bedoop application 36 designed to service the discerned Bedoop data.

Sometimes the system 28 may encounter a Bedoop object for which several different responses may be appropriate. In the case of a printed office document, for example, one response may be as described above – to present the electronic version of the file on a computer, ready for editing. But other responses may also be desired, such as writing an email message to the author of the printed document, with the author's email address already specified in the message address field, etc.

Such different responses may be handled by different Bedoop applications, or may be options that are both provided by a single Bedoop application. In the former case, when the CLASS/DNS IDs are decoded and provided to the operating system, the registry indicates that

there are two (or more) programs that might be invoked. The operating system can then present a dialog box to the user inviting the user to specify which form of response is desired. Optionally, a default choice can be made if the user doesn't specify within a brief period (e.g., three seconds). The operating system can then launch the Bedoop application corresponding to the chosen response.

A similar arrangement can be employed if a single Bedoop application can provide both responses. In such case the operating system launches the single Bedoop application (since there is no ambiguity to be resolved), and the application presents the choice to the user. Again, the user can select, or a default choice can be automatically made.

In the just-described situations, the user can effect the choice by using the keyboard or mouse – as with traditional dialog boxes. But Bedoop provides another, usually easier, form of interaction. The user can make the selection through the optical sensor input. For example, moving the object to the right can cause a UI button on the right side of the dialog box to be selected; moving the object to the left can cause a UI button on the left side of the dialog box to be selected; moving the object towards the camera can cause the selected button to be activated. Many other such techniques are possible, as discussed below.

If the registry 34 does not recognize, or otherwise does not know how to respond to Bedoop data of that particular CLASS/DNS, the registry launches a default Bedoop client application. This client application, in turn, directs a web browser 40 on the local Bedoop system 28 to communicate with a remote master registration server computer 42. The local computer forwards the Bedoop data to this master server. The master server 42 examines the CLASS ID, and forwards the Bedoop data (directly, or through intervening servers) to a corresponding CLASS server 44. (A single server may handle Bedoop data of several classes, but more typically there is a dedicated server for each CLASS.)

Each CLASS server 44 serves as the root of a tree 46 of distributed DNS servers. A DNS server 48a, for example, in a first tier 50 of the DNS server tree, may handle Bedoop data having DNS IDs beginning with "000." Likewise, DNS server 48b may handle Bedoop data having DNS IDs beginning with "001," etc., etc.

Each DNS server in the first tier 50 may, in turn, route Bedoop data to one of 8 servers in a second tier of the tree, in accordance with the fourth- through sixth bits of the DNS data. The tree continues in this fashion until a terminal level of DNS leaf node servers 56.

Ultimately, Bedoop data routed into this network reaches a DNS leaf node server 56. That leaf node server may handle the Bedoop data, or may redirect the local Bedoop system to a further server 58 that does so. That ultimate server – whether a DNS leaf node server or a further server – can query the local Bedoop system for further information, if necessary, and can either instruct the local Bedoop system how to respond, or can undertake some or all of the response itself and simply relay appropriate data back to the local Bedoop system.

In arrangements in which the local Bedoop system is redirected, by the DNS leaf node server, to a further server that actually handles the response, access to the further server may be through a port 59 (e.g., a special URL) tailored to receipt of Bedoop data.

In a typical implementation, most or all of the servers are mirrored, or otherwise replicated/redundant, so that failure of individual computers does not impair operation of the system.

Caching can be provided throughout the trees of servers to speed responses. That is, responses by leaf nodes for certainly commonly-encountered CLASS/DNS IDs can be temporarily stored earlier in the tree(s). Bedoop data, propagating through the server network, can prompt a response from an intermediate server if there is a cache hit.

If desired, Bedoop traffic through the above-detailed server trees can be monitored to collect demographic and statistical information as to what systems are sending what Bedoop data, etc. One use of such information is to dynamically reconfigure the DNS network to better balance server loads, to virtually relocate DNS resources nearer regions of heavy usage, etc. Another use of such information is for marketing purposes, e.g., to promote certain Bedoop features and applications within user groups (e.g., internet domains) that seem to under-utilize those features.

Within certain user networks that are linked to the internet, e.g., corporate networks, Bedoop data that isn't handled within the originating Bedoop system may first be routed to a Bedoop name server within the corporate network. That server will recognize certain types of Bedoop data, and know of resources within the corporate network suitable for handling same. Referral to such resources within the corporate network will be made, where possible. These resources (e.g., corporate servers) may respond to Bedoop data in a way customized to the corporate preferences. If the corporate Bedoop name server does not know of a resource within the corporate network that can respond to the Bedoop data, the corporate name server

then routes the data to the public Bedoop network described above. (Such referral can be to the master registration server or, to the extent the corporate name server knows the addresses of appropriate servers within the DNS server tree, or of the further servers to which DNS servers may point for certain Bedoop data, it can redirect the local Bedoop system accordingly.)

In typical rich Bedoop implementations, local systems may have libraries of Bedoop services, applications, or protocols. Some may be unique to that computer. Others may be commonly available on all computers. Some may be highly secure, employing encryption and/or anti-hacking measures, or data protocols that are not generally recognized. Others may be shareware, or the result of open-source programming efforts.

Greeting Cards, Birthday Cards, Etc.

In accordance with a further embodiment of the invention, greeting cards and the like are encoded (e.g., by texturing, printing, etc.) with Bedoop data. On receiving such a card, a recipient holds it in front of the image capture device on a laptop or other computer. The computer responds by displaying an internet web page that has a stock- or customized- presentation (image, video, audio-video, etc.) to complement that presented on the greeting card.

The web site presentation can be personalized by the sender (e.g., with a text message, recent family photographs, etc.), either at the point of card sale, or sometime after the card is purchased. In the latter case, for example, the card can be serialized. After taking the card home, the purchaser can visit the card vendor's web site and enter the card serial number in an appropriate user interface. The purchaser is then presented with a variety of simple editing tools to facilitate customization of the web greeting. When the sender is finished designing the web greeting, the finished web page data is stored (by software at the vendor's web site) at a site corresponding to the serial number.

When the card is received by a recipient and held in front of a Bedoop sensor, CLASS, DNS, and UID data is decoded from the card. The CLASS and DNS data are used to navigate the earlier-described server network to reach a corresponding DNS leaf node server (perhaps maintained by the Hallmark greeting card company). That leaf node server indexes a table, database, or other data structure with the UID from the Bedoop data, and obtains from that

data structure the address of an ultimate web site – the same address at which the web greeting customized by the sender was stored. That address is provided by the DNS leaf node server back to the local computer, with instructions that the web page at that address be loaded and displayed (e.g., by HTML redirection). The local computer complies, presenting the customized web greeting to the card recipient.

In the just-described embodiment, in which a pre-encoded card is purchased by a sender and the web-display is then customized, the address of the web site is typically determined by the card vendor. But this need not be the case. Likewise, the card need not be “purchased” in the typical, card-shop fashion.

To illustrate the foregoing alternatives, consider the on-line acquisition of a greeting card, e.g., by visiting a web site specializing in greeting cards. With suitable user-selection (and, optionally, customization), the desired card can be printed using an ink-jet or other printer at the sender’s home. In such case, the Bedoop data on the card can be similarly customized. Instead of leading to a site determined by the card vendor, the data can lead to the sender’s personal web page, or to another arbitrary web address.

To effect such an arrangement, the sender must arrange for a DNS leaf node server to respond to a particular set of Bedoop data by pointing to the desired web page. While individuals typically will not own DNS servers, internet service providers commonly will. Just as AOL provides simple tools permitting its subscribers to manage their own modest web pages, internet service providers can likewise provide simple tools permitting subscribers to make use of DNS leaf node servers. Each subscriber may be assigned up to 20 UIDs (under a particular CLASS and DNS). The tools would permit the users to define a corresponding web address for each UID. Whenever a Bedoop application led to that DNS leaf node server, and presented one of those UIDs, the server would instruct the originating computer to load and present the web page at the corresponding web address.

Prior to customizing the greeting card, the sender uses the tool provided by the internet service provider to store the address of a desired destination web address in correspondence with one of the sender’s available UIDs. When customizing the greeting card, the sender specifies the Bedoop data that is to be encoded, including the just-referenced UID. The greeting card application encodes this data into the artwork and prints the resulting card.

When this card is later presented to a Bedoop system by the recipient, the recipient's system loads and displays the web page specified by the sender.

Commerce in Bedoop Resources

5 In the just-described arrangement, internet service providers make available to each subscriber a limited number of UUIDs on a DNS server maintained by the service. Business enterprises typically need greater Bedoop resources, such as their own DNS IDs (or even their own CLASS ID(s)).

10 While variants of the Bedoop system are extensible to provide an essentially unlimited number of CLASS IDs and DNS IDs, in the illustrated system these resources are limited. Public service, non-profit, and academic applications should have relatively generous access to Bedoop resources, either without charge or for only a modest charge. Business enterprises, in contrast, would be expected to pay fees to moderate their potentially insatiable demand for the resources. Small businesses could lease blocks of UUIDs under a given CLASS/DNS ID.

15 Larger businesses could acquire rights to entire DNS IDs, or to entire CLASS IDs (at commensurately greater fees).

 Web-based systems for assigning DNS IDs (and CLASS IDs) can be modeled after those successfully used by Internic.com, and now Networksolutions.com, for registration of internet domains. The user fills out a web-based form with names, addresses, and billing

20 information; the system makes the necessary changes to all of the hidden system infrastructure – updating databases, routing tables, etc., in servers around the world.

Controlled-Access ID

25 Just as the above-described embodiment employed an ink-jet printer to produce a customized-Bedoop greeting card, the same principles can likewise be applied to access-control objects, such as photo-IDs.

 Consider an employment candidate who will be interviewing at a new employer. The candidate's visit is expected, but she is not recognized by the building's security personnel. In this, and many other applications, arrangements like the following can be used:

The employer e-mails or otherwise sends the candidate an access code. (The code can be encrypted for transmission.) The code is valid only for a certain time period on a given date (e.g., 9:00 a.m. – 11:00 a.m. on June 29, 1999).

Upon receipt of the access code, the candidate downloads from the web site of the state Department of Motor Vehicles the latest copy of her driver's license photo. The DMV has already encoded this photo with Bedoop data. This data leads to a state-run DNS leaf node server 56. When that server is presented with a UID decoded from a photograph, the server accesses a database and returns to the inquiring computer a text string indicating the name of the person depicted by the photograph.

The candidate incorporates this photo into an access badge. Using a software application (which may be provided especially for such purposes, e.g., as part of an office productivity suite), the photo is dragged into an access badge template. The access code emailed from the employer is also provided to this application. On selecting "Print," an ink-jet printer associated with the candidate's computer prints out an access badge that includes her DMV photo and her name, and is also steganographically encoded in accordance with the employer-provided access code.

The name printed on the badge is obtained (by the candidate's computer) from the DMV's DNS server, in response to Bedoop data extracted from the photograph. (In this application, unlike most, the photograph is not scanned as part of a Bedoop process. Instead, the photograph is already available in digital form, so the Bedoop decoding proceeds directly from the digital representation.)

For security purposes, the access code is not embedded using standard Bedoop techniques. Instead, a non-standard format (typically steganographic) is employed. The embedding of this access code can span the entire face of the card, or can be limited to certain regions (e.g., excluding the region occupied by the photograph).

On the appointed day the candidate presents herself at the employer's building. At the exterior door lock, the candidate presents the badge to an optical sensor device, which reads the embedded building access code, checks it for authenticity and, if the candidate arrived within the permitted hours, unlocks the door.

Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the

candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e.g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)

The guard checks the name returned by the Bedoop system with the name printed on the badge. On seeing that the printed and Bedoop-decoded names match (and optionally checking the door log to see that a person of that name was authorized to enter and did so), the security guard can let the candidate pass.

It will be recognized that the just-described arrangement offers very high security, yet this security is achieved without the candidate ever previously visiting the employer, without the employer knowing what the candidate looks like, and by use of an access badge produced by the candidate herself.

Variants of such home-printed badge embodiments find numerous applications. Consider purchasing movie- or event-tickets over the web. The user can print an access ticket that has an entry code embedded therein. On arriving at the theater or event, the user presents the ticket to an optical scanning device, which decodes the entry code, checks the validity of same, authorizes the entry, and marks that entry code as having been used (preventing multiple uses of tickets printed with the same code).

Ink-Jet Printing

In the foregoing discussions, reference has been made to use of ink-jet printing as a means for providing steganographically encoded indicia on substrates. The following discussion expands on some of the operative principles.

The basic physics and very low level analog electronic operation of ink-jet printers (sometimes termed bubble-jet printers) are ideally suited to support very-light-tint background digital watermarking on any form of substrate. (Watermarking through apparent "tinting" of substrates is discussed in copending application 09/127,502.) In general, the statement, "if you can print it with an ink jet printer, you can watermark it" is largely accurate, even for (perhaps especially for) simple text documents. Indeed, there is a degree of flexibility and

control in the ink-jet printing realm that is not as generally available in more traditional printing technologies, such as commercial offset printing and other plate-based technologies. (This is not to say that ink-jet has better quality than plate-based technologies; it has more to do with the statistics of ink droplets than anything else.) Heavier tint backgrounds are possible as well, where the continuum ranges from very light background tinting, where the casual observer will see “white paper,” all the way through heavily inked patterned backgrounds, and photographs themselves, and everything in between.

In some embodiments, the ink-jet driver software is modified to provide lower-level control of individual droplet emission than is provided in existing printer drivers, which are naturally optimized for text and graphics. In some such embodiments, the “watermarking” print mode is another option from which the user can select (e.g., in addition to High Quality, Econo-Fast, etc.), or the selection can be made automatically by application software that is printing watermarked data.

In more sophisticated embodiments, the watermark data is applied to the printer driver software independently of the other image/text data. The printer driver is arranged to eject droplets in the usual print density for the image/text data, and at a more accurately-controlled, finer density for the separately-applied watermark data. (The latter may be effected as a slight modulation signal on the former.) This arrangement provides for essentially transparent integration into existing printer environments – no one need worry about the watermarking capability except the software applications that specifically make use of same.

Consumer Marking of Web-Based Materials

Various items of printed media can originate off the web, yet be printed at home. Examples include movie tickets, coupons, car brochures, etc. Bedoop data can be added, or modified, by the software application or by the printer driver at the time of printing. (Alternatively, the Bedoop data can be customized to correspond to the user before being downloaded to the user’s system for printing.)

One advantage to Bedoop-encoding printed images locally, as opposed to Bedoop-encoding the image files prior to downloading for local printing, is that the encoding can be tailored in accordance with the particular properties of the local printer (e.g., to increase robustness or decrease visibility) – properties not generally known to a remote server.

In one particular example, the UID field in the Bedoop data can be written with a value that serves as an index to a database of user profiles, permitting later systems to which the printed item is presented to personalize their response in accordance with the profile data.

In another example, the UID field serves an authentication purpose, e.g., to verify that the printed medium actually was printed at a particular place, or by a particular user or at a particular time.

Coffee Mug

At retail coffee outlets, customers commonly order the same drink day after day (“half-decaf, short, skinny latte”). Some customers present personal coffee mugs to the cashier, preferring the sensation of ceramic or metal to paper, and avoiding the trash/recycle dilemma.

The drinker’s “regular” order can be Bedoop-encoded either on the mug itself or, more commonly, on an adhesive label applied to the mug. The encoding can be in addition to other aesthetic imagery (e.g., artwork or a photo), or the marking can be purely data. Labels the size of postage stamps may be used.

On handing the mug to the cashier, the customer can simply say “the regular.” The cashier passes the mug in front of the optical scanning device of a Bedoop system associated with the cash register. The system steganographically decodes the data and provides the corresponding order (“half-decaf, short, skinny latte”), either textually or audibly (e.g., by a voice synthesizer) to the cashier or the barrista. The cash register system also knows the current price of the requested drink, and rings up the charge accordingly.

Labels of the type described can be available to the cashier on pre-printed rolls, just as with other adhesive stickers, or can be printed on-demand. (Small label printers may be best suited in the latter case, given space constraints in retail outlets.) Customers ordering drinks for personal mugs may be invited to take a label corresponding to their just-ordered drink and apply it to their mug for future use.

In variants on this basic theme, the mug label can be further encoded (or a supplemental label can be provided and encoded) with electronic payment information, such as the customer’s credit card number, or the number of a debit account maintained by the coffee merchant for that customer. When the mug is scanned for the drink order, the system

likewise detects the payment information and charges the corresponding fee to the appropriate account. (For security reasons, the system may be arranged so that the mug cannot be used to authorize more than, say \$5 of coffee drink purchases per day.)

5 In another variant on this theme, the system maintains an electronic log of coffee purchases made by the customer and, in accordance with then-prevailing marketing considerations, rewards the customer with a free drink after 8 or 12, etc., drinks have been purchased.

10 In still another variant on this theme, regular customers who use Bedoop-labeled mugs can participate in periodic promotions in which, for example, every N^{th} such customer is rewarded with a cash or merchandise prize. Bells go off when the N^{th} mug is scanned. (N can be a fixed number, such as 500, or can be a random number – typically within a known range or with a known mean.)

Smart Elevators

15 In accordance with another embodiment of the invention, a building elevator is provided with one or more optical capture devices. Each device examines monitors the contents of the elevator chamber looking for Bedoop encoded objects, such as ID badges.

20 On sensing a Bedoop-encoded object, the elevator can determine – among other data – the floor on which the wearer's office is located. The system can then automatically direct the elevator to that floor, without the need for the person to operate any buttons. (The elevator's button panel can be provided with a new, override button that can be operated to un-select the most recently selected floor(s), e.g., in case a user wants to travel to a different floor.)

25 To aid in identification, the Bedoop objects (e.g., badges) can be colored a distinctive color, permitting the system to more easily identify candidate objects from other items within the optical capture devices' field of view. Or the object can be provided with a retro-reflective coating, and the elevator can be equipped with one or more illumination sources of known spectral or temporal quality (e.g., constant infra red, or constant illumination with a single- or multi-line spectrum, or a pulsed light source of known periodicity; LEDs or semiconductor lasers, each with an associated diffuser, can be used for each the foregoing and can be paired
30 with the image capture devices). Other such tell-tale clues can likewise be used to aid in object location. In all such cases, the optical capture device can sense the tell-tale clue(s)

using a wide field of view sensor. The device can then be physically or electronically steered, and/or zoomed, to acquire a higher resolution image of the digitally-encoded object suitable for decoding.

5 Magazines

Magazine (and newspaper) pages can be steganographically encoded with Bedoop data to provide another “paper as portal” experience. As with the earlier described office document case, the encoded data yields an address to a computer location (e.g., a web page) having the same, or related, content.

10 In one exemplary embodiment, the blank magazine page stock is Bedoop-encoded prior to printing. The watermarking can be performed by high speed ink-jet devices, which splatter a fine pattern of essentially imperceptible ink droplets across each page. Each page can be differently watermarked so that, on decoding, page 21 of a magazine can be distinguished from page 22 of the same magazine (and page 106 of the June 21, 1999, issue
15 can be distinguished from page 106 of the June 28, 1999, issue). If desired, each page can be further segregated into regions – either in accordance with the actual boundaries of articles that will later be printed on the pages, or in a grid pattern, e.g., of 3 columns across by 5 rows high. Each region conveys a distinct Bedoop code, permitting different portions of the page to lead to different web data.)

20 After watermarking and printing, the pages thus produced are bound in the usual fashion with others to form the finished magazine. (Not all pages in the magazine need to be watermarked.)

Of course, the watermarking can be effected by processes other than ink-jet printing. For example, texturing by pressure rollers is another option well suited for the large volumes
25 of paper to be processed.

On presenting a magazine to the optical scanner device of a Bedoop-compliant computer, the computer senses the Bedoop data, decodes same, and launches a web browser to an internet address corresponding to the Bedoop data. If the magazine page is an advertisement, the internet address can provide information complementary to the
30 advertisement. For example, if the magazine page is an advertisement for a grocery item, the Bedoop data can identify a web page on which recipes using the advertised item are

presented. If the magazine page includes a photo of a tropical beach, the Bedoop data can lead to a travel web page (e.g., hosted by Expedia or other travel service) that presents fare and lodging information useful to a reader who wants to vacation at the illustrated beach. (The fare information can be customized to the reader's home airport by reference to user profile data stored on the user's computer and relayed to the web site to permit customization of the displayed page.)

The data to which the Bedoop data leads needn't be static; it can be updated on a weekly, daily, or other basis. Thus, if a months-old magazine page is presented to a Bedoop device, the resultant data can be up-to-the-minute.

In the case of advertising, the inclusion of Bedoop data increases the value of the ad to the advertiser, and so merits a higher charge to the advertiser from the magazine publisher. This higher charge may be shared with the enterprise(s) that provides the Bedoop technology and infrastructure through which the higher value is achieved.

Business Card Applications

Conventional business cards can be steganographically encoded with Bedoop data, e.g., by texturing, watermark tinting, ink-jet splattering, text steganography, etc. As with many of the earlier-described embodiments, the steganographic encoding is tailored to facilitate decoding in the presence of arbitrary rotation or scale distortion of the card introduced during scanning. (Some such techniques are shown, e.g., in applicant's related patents identified above. Various other techniques are known to artisans.)

When a recipient of a business card holds it in front of a Bedoop sensor, the operating system on the local system launches a local Bedoop application. That local Bedoop application, in turn, establishes an external internet connection to a remote business card server. The address of that server may already be known to the local Bedoop application (e.g., having been stored from previous use), or the local Bedoop system can traverse the above-described public network of DNS servers to reach the business card server.

A database on the business card name server maintains a large collection of business card data, one database record per UID. When that server receives Bedoop data from a local Bedoop system, it parses out the UID and accesses the corresponding database record. This record typically includes more information than is commonly printed on conventional

business cards. Sample fields from the record may include, for example, name, title, office phone, office fax, home phone, home fax, cellular phone, email address, company name, corporate web page address, personal web page address, secretary's name, spouse's name, and birthday. This record is transmitted back to the originating Bedoop system.

5 The local Bedoop system now has the data, but needs further instruction from the user as to how it should be processed. Should a telephone number be dialed? Should the information be entered into a personal contact manager database (e.g., Outlook) on the local system? Etc.

10 In an exemplary embodiment, the local system presents the available choices to the user, e.g., by textual prompts, synthesized voice, etc. The user responds by manipulating the business card in a manner prompted by the system (e.g., move down to telephone at office; move up to telephone at home; move right to access corporate web page; move left to access personal web page; rotate left to enter certain elements from the database record (filtered in accordance with a template) into personal contact manager database, etc. The local Bedoop system responds accordingly.

15 Some card givers may choose to make additional information available to card recipients -- information beyond that known in prior art contact-management software applications. For example, one of the choices presented by a local Bedoop system in response to presentation of a business card may be to review the card-giver's personal calendar. (The card-giver can maintain his or her personal calendar on a web-accessible computer.) By such arrangement, the card-recipient can learn when the card-giver may be found in the office, when appointments might be scheduled, etc., etc.

20 Typically, access to this web-calendar is not available to casual web browsers, but is accessible only in response to Bedoop data (which may thus be regarded as a form of authentication or password data).

25 Some users may carry several differently-encoded cards, each with a different level of access authorization (e.g., with different UIDs). Thus, some cards may access a biographical page without any calendar information, other cards may access the same or different page with access enabled to today's calendar, or this week's calendar, only, and still other cards (e.g., the "spouse" card) may access the same or different page with access enabled for the card-giver's complete calendar. The user can distribute these different cards to different

persons in accordance with the amount of personal information desired to be shared with each.

In accordance with a related embodiment, the database record corresponding to Bedoop business card data can include a “now” telephone number field. This field can be continually-updated throughout the day with the then-most-suitable communications channel to the card-giver. When the card-giver leaves home to go to the office, or leaves the office for a trip in the car, or works a week at a corporate office in another town, etc., this data field can be updated accordingly. (A pocket GPS receiver, with a wireless uplink, can be carried by the person to aid in switching the “now” number among various known possibilities depending on the person’s instantaneous position.) When this database record is polled for the “now” number, it provides the then-current information.

Consider a Bedoop-enabled public telephone. To dial the phone, a business card is held in front of the Bedoop sensor (or slid through an optical scanner track). The phone interrogates the database at the business card server for the “now” number and dials that number.

To update the any of the fields stored in the database record, the card giver can use a special card that provides write-authorization privileges. This special card can be a specially encoded version of the business card, or can be another object unique to the card-giver (e.g., the card-giver’s driver’s license).

The reference to business cards and personal calendars is illustrative only. Going back a century, “calling cards” were used by persons whose interests were strictly social, rather than business. The just-discussed principles can be similarly applied. Teenagers can carry small cards to exchange with new acquaintances to grant access to private dossiers of personal information, favorite music, artwork, video clips, etc. The cards can be decorated with art or other indicia that can serve purposes wholly unrelated to the Bedoop data steganographically encoded therein.

Gestural Input

A Bedoop system can determine the scale state, rotation state, X-Y offset, and differential scale state, of an object by reference to embedded calibration data, or other techniques. If the scan device operates at a suitably high frame rate (e.g., five or ten frames

per second), change(s) in any or all of these four variables can be tracked over time, and can serve as additional input.

In an earlier-discussed example, moving an object to the left or right in front of the Bedoop scanner caused a left- or right-positioned button in a dialog box to be selected. This is a change in the X-Y offset of the scanned object. In that earlier example, moving the object inwardly towards the camera caused the selected button to be activated. This is a change in the scale state of the scanned object.

In similar fashion, twisting the object to the left or right can prompt one of two further responses in a suitably programmed Bedoop application. (This is a change in the rotation state.) Likewise, tilting the object so that one part is moved towards or away from the camera can prompt one of two further responses in the application. (This is a change in the differential scale state.)

In the business card case just-discussed, for example, the card can be held in front of the Bedoop scanner of a computer. If the card is twisted to the left, the computer opens a web browser to a web page address corresponding to Bedoop data on the card. If the card is twisted to the right, the computer opens an e-mail template, pre-addressed to an e-mail address indicated by the card.

In other examples, twisting an object to move the right edge towards the scanner can be used to effect a right mouse click input, and twisting the object to move the right edge away from the scanner can be used to effect a left mouse click input.

Simultaneous changes in two of these four positioning variables can be used to provide one of four different inputs to the computer (e.g., (a) twisting left while moving in; (b) twisting left while moving out; (c) twisting right while moving in; and (d) twisting right while moving out). Simultaneous changes to three or all four of these variables can similarly be used to provide one of eight or sixteen different inputs to the computer.

Simultaneous manipulations of the object in two or more of these modes is generally unwieldy, and loses the simple, intuitive, feel that characterizes manipulation of the object in one mode. However, a similar effect can be achieved by sequential, rather than simultaneous, manipulation of the card in different modes (e.g., twist left, then move in). Moreover, sequential manipulations permit the same mode to be used twice in succession (e.g., move in,

then move out). By such sequential manipulations of the object, arbitrarily complex input can be conveyed to the Bedoop system.

(It will be recognized that a digitally-encoded object is not necessary to the gestural-input applications described above. Any object (talisman) that can be distinguished in the image data can be manipulated by a user in the manners described above, and an appropriate system can recognize the movement of the object and respond accordingly. The provision of digital data on the object provides a further dimension of functionality (e.g., permitting the same gesture to mean different things, depending on the digital encoding of the object being manipulated), but this is not essential.

Moreover, even within the realm of digitally-encoded gestural talismans, steganographic encoding is not essential. Any other known form of optically-recognizable digital encoding (e.g., 1D and 2D bar codes, etc.) can readily be employed.

In an illustrative embodiment, a business card or photograph is used as the talisman, but the range of possible talismans is essentially unlimited.

Gestural Decoding Module

There are various ways in which the Bedoop system's decoding of gestural input can be effected. In some Bedoop systems, this functionality is provided as part of the Bedoop applications. Generally, however, the applications must be provided with the raw frame data in order to discern the gestural movements. Since this functionality is typically utilized by many Bedoop applications, it is generally preferable to provide a single set of gestural interpretation software functions (commonly at the operating system level) to analyze the frame data, and make available gestural output data in standardized form to all Bedoop applications.

In one such system, a gestural decoding module tracks the encoded object within the series of image data frames, and outputs various parameters characterizing the object's position and manipulation over time. Two of these parameters indicate the X-Y position of the object within current frame of image data. The module can identify a reference point (or several) on the object, and output two corresponding position data (X and Y). The first represents the horizontal offset of the reference point from the center of the image frame, represented as a percentage of frame width. A two's complement representation, or other

representation capable of expressing both positive and negative values, can be used so that this parameter has a positive value if the reference point is right of center-frame, and has a negative value if the reference point is left of center frame. The second parameter, Y, similarly characterizes the position of the reference point above or below center-frame (with above-being represented by a positive value). Each of these two parameters can be expressed
5 as a seven-bit byte. A new pair of X, Y parameters is output from the gestural decoding module each time a new frame of image data is processed.

In many applications, the absolute X-Y position of the object is not important. Rather, it is the movement of the object in X and Y from frame-to-frame that controls some aspect of
10 the system's response. The Bedoop application can monitor the change in the two above-described parameters, frame to frame, to discern such movement. More commonly, however, the gestural decoding module performs this function and outputs two further parameters, X' and Y'. The former indicates the movement of the reference point in right/left directions since the last image frame, as a percentage of the full-frame width. Again, this parameter is
15 represented in two's complement form, with positive values representing movement in the rightward direction, and negative values representing movement in the leftward direction. The later parameter similarly indicates the movement of the reference point in up/down directions since the last frame.

The scale, differential scale, and rotation states of the object can be similarly analyzed
20 and represented by parameters output from the gestural decoding module.

Scale state can be discerned by reference to two (or more) reference points on the object (e.g., diagonal corners of a card). The distance between the two points (or the area circumscribed by three or more points) is discerned, and expressed as a percentage of the diagonal size of the image frame (or its area). A single output parameter, A, which may be a
25 seven-bit binary representation, is output.

As with X-Y data, the gestural decoding module can likewise monitor changes in the scale state parameter since the last frame, and product a corresponding output parameter A'. This parameter can be expressed in two's complement form, with positive values indicating movement of the object towards the sensor since the last frame, and negative values indicating
30 movement away.

A differential scale parameter, B, can be discerned by reference to four reference points on the object (e.g., center points on the four edges of a card). The two points on the side edges of the card define a horizontal line; the two points on the top and bottom edges of the card define a vertical line. The ratio of the two line lengths is a measure of differential scale. This ratio can be expressed as the shorter line's length as a percentage of the longer line's length (i.e., the ratio is always between zero and one). Again, a two's complement seven-bit representation can be used, with positive values indicating that the vertical line is shorter, and negative values indicating that the horizontal line is shorter. (As before, a dynamic parameter B' can also be discerned to express the change in the differential scale parameter B since the last frame, again in two's complement, seven bit form.)

A rotation state parameter C can be discerned by the angular orientation of a line defined by two reference points on the object (e.g., center points on the two side edges of a card). This parameter can be encoded as a seven-bit binary value representing the percentage of rotational offset in a clockwise direction from a reference orientation (e.g., horizontal). (The two reference points must be distinguishable from each other regardless of angular position of the object, if data in the full range of 0 – 360 degrees is to be represented. If these two points are not distinguishable, it may only be possible to represent data in the range of 0-180 degrees.) As before, a dynamic parameter C' can also be discerned to express the change in the rotation state parameter C since the last frame. This parameter can be in seven bit, two's complement form, with positive values indicating change in a clockwise rotation

The foregoing analysis techniques, and representation metrics, are of course illustrative only. The artisan will recognize many other arrangements that can meet the needs of the particular Bedoop applications being served.

In the illustrative system, the Bedoop application programs communicate with the gestural decoding module through a standardized set of interface protocols, such as APIs. One API can query the gestural input module for some or all of the current position parameters (e.g., any or all of X, Y, A, B, and C). The module responds to the calling application with the requested parameter(s). Another API can query the gestural input module for some or all of the current movement data (e.g., any or all of X', Y', A', B' and C'). Still another API can request the gestural decoding module to provide updated values for some or all of the position or movement data on a running basis, as soon as they are discerned

from each frame. A complementary API discontinues the foregoing operation. By such arrangement, all of the gestural data is available, but the Bedoop application programs only obtain the particular data they need, and only when they ask for it.

In Bedoop applications that communicate with external servers, just the Bedoop data (i.e., CLASS, DNS, and optionally UID) may initially be sent. If the remote server needs to consider gestural data in deciding how to respond, the remote server can poll the local Bedoop system for the necessary data. The requested gestural data is then sent by the local Bedoop system to the remote server in one or more separate transmissions.

In other embodiments, since the gestural data is of such low bandwidth (e.g., roughly 56 bits per image frame), it may routinely and automatically be sent to the remote computer, so that the gesture data is immediately available in case it is needed. In an illustrative implementation, this data is assembled into an 8-byte packet, with the first byte of the packet (e.g., the X parameter) being prefixed with a “1” sync bit, and subsequent bytes of the packet being prefixed with “0” sync bits. (The sync bits can be used to aid in accurate packet decoding.)

In some embodiments, it is useful to provide for an extension to the normal 64-bit Bedoop length to accommodate an associated packet of gestural data. This can be effected by use of a reserved bit, e.g., in the UID field of the Bedoop packet. This bit normally has a “0” value. If it has a “1” value, that indicates that the Bedoop data isn’t just the usual 64 bits, but instead is 128 bits, with the latter 64 bits comprising a packet of gestural data.

Similar extension protocols can be used to associate other ancillary data with Bedoop data. A different reserved bit in the UID field, for example, may signal that a further data field of 256 bits follows the Bedoop data – a data field that will be interpreted by the remote computer that ultimately services the Bedoop data in a known manner. (Such bits may convey, e.g., profile data, credit card data, etc.) The appended data field, in turn, may include one or more bits signaling the presence of still further appended data.

Grandmothers

It is a common complaint that computers are too complex for most people. Attempts to simplify computer-user interaction to facilitate use by less experienced users usually serve to frustrate more experienced users.

In accordance with another embodiment of the present invention, the sophistication of a computer user is steganographically indicated on a talisman used by that user to interact with the system. The computer detects this steganographically-encoded data, and alters its mode of interacting with the user accordingly.

5 Consider internet browser software. Experienced users are familiar with the different functionality that can be accessed, e.g., by various drop-down menus/sub-menus, by the keyboard shortcuts, by the menus available via right-clicking on the mouse, by manipulating the roller mouse scroll wheel and scroll button, etc., etc. Grandmothers of such users, typically, are not so familiar.

10 Although gestural interfaces hold great promise for simplifying user-computer interaction, the same dichotomy between experienced users and inexperienced users is likely to persist, frustrating one class of user or the other.

To help close this gap, a computer system according to this embodiment of the invention responds to gestures in different manners, depending on the expertise level indicated by encoding of the talisman. For an expert user, for example, the gestural interface active in the internet browser software may display the stored list of Favorite web addresses in response to tipping the left edge of the talisman towards the optical sensor. Once this list is displayed, the expert user may rotate the talisman to the right to cause the highlighting to scroll down the list from the top. Rotating the talisman to the left may scroll the list of Favorites up from the bottom. The speed of scrolling can be varied in accordance with the degree of rotation of the talisman from a default orientation.

15 In contrast, for the novice user, these talisman manipulations may be confounding rather than empowering. Tipping the left edge of the talisman towards the sensor may occur as often by mistake as on purpose. For such users, a more satisfactory interface may be provided by relying on simple X-Y movement of the talisman to move an on-screen cursor, with a movement of the talisman towards the sensor to serve as a selection signal (i.e., like a left-mouse click).

25 (In the example just-cited, the expert user summoned a list of Favorite web sites. Different "Favorites" lists can be maintained by the computer – each in association with different talismans. A husband who uses one talisman is provided a different "Favorites" list than a wife who uses a different talisman.)

Printed Pictures

In accordance with this aspect of the invention, a printed photograph can be steganographically encoded with Bedoop data leading to information relating to the depicted person (e.g., contact information, biographical information, etc.).

Such a photograph can be presented to a Bedoop sensor on a telephone. In a simple embodiment, the telephone simply processes the Bedoop data to obtain a corresponding default telephone number, and dials the number. In other embodiments, various options are possible, e.g., dial home number or dial work number. On presenting the photograph to the telephone, for example, moving the photo to the left may dial the person at home, while moving the photo to the right may dial the person at work.

As telephones evolve into more capable, multi-function devices, other manipulations can invoke other actions. In a computer/telephone hybrid device, for example, rotating the photo counterclockwise may launch a web browser to an address at which video data from a web cam at the pictured person's home is presented. Rotating the photo clockwise may present an e-mail form, pre-addressed to the e-mail address of the depicted person. Moving the photo to the right may query a database on the system for other photographs depicting the same individual or subject, which can be presented in response to further user input. Etc.

In this and other embodiments, it is helpful for the Bedoop device to prompt the user to aid in manipulating the object. This can be done audibly (e.g., "move photo left to dial at home") or by visual clues (e.g., presenting left- or right-pointing arrows).

Bedoop data in photographs can also be used to annotate the photographs, as with notes on the back of a photograph, or printed under the photograph in a photo album. The Bedoop data can lead to a remote database, where the photograph owner is permitted to enter a textual (or audio) narrative in association with each photograph's UID. Years later, when some of the names have been forgotten, the photograph can be positioned in front of a Bedoop sensor, and the system responds by providing the annotation provided by the photograph owner years earlier.

Drivers Licenses and Other Cards

Drivers licenses, social security cards, or other identity documents may be encoded by the issuing authority with Bedoop data that permits access to the holder's personal records over the web. On presenting the document to a Bedoop system, the system directs a web browser to a private address corresponding to data encoded on the document. At that address, the holder of the document can review governmental records, such as state or federal tax return data, social security entitlements, etc., as well as privately-maintained records, such as credit records, etc. User selection among various functions can be effected by spatial manipulation of the document. (Entry of additional data, such as social security number or mother's maiden name, may be required of the user to assure privacy in case the document is lost or stolen.)

By manipulating a driver's license in front of a Bedoop sensor, a user can request renewal of the driver's license, and authorize payment of the corresponding fee.

Bank cards (debit, credit, etc.) can similarly be encoded with Bedoop data to permit the holder to access bank records corresponding to the bank card account. (Entry of a PIN code may be required to assure privacy.)

Such documents can also be used to access other personal data. One example is e-mail. A traveler might pause at a Bedoop kiosk at an airport and present a driver's license. Without anything more, the kiosk may present email that is waiting for the traveler on an associated display screen.

On recognizing a driver's license, the kiosk can access a remote site (which may be maintained by the Department of Motor vehicles, another government entity, a private entity, or by the traveler), authenticating the operation by presenting Bedoop data encoded on the license, and obtaining information that the person has pre-approved for release in response to such authorized access. This information can include e-mail account and password information. Using this information, the kiosk queries the corresponding e-mail server, and downloads a copy of recently received mail for presentation at the kiosk. (A user-entered PIN number may be required at some point in the process, e.g., in querying the remote site for sensitive e-mail password data, before presenting the downloaded e-mail for viewing, etc., to ensure privacy.)

Other cards carried in wallets and purses can also be encoded to enable various functions. The local sandwich shop that rewards regular customers by awarding a free sandwich after a dozen have been purchased can encode their frequent-buyer card with Bedoop data leading to the shop's web-based sandwich delivery service. Or the frequent-buyer card can be eliminated, and customers can instead wave their business card or other identity document in front of the shop's Bedoop sensor to get purchase credit in a tally maintained by the sandwich shop's computer.

Food stamps, health insurance cards, and written medical prescriptions, can likewise be encoded with digital data to enable the provision of new functionality.

At large trade shows, such as COMDEX, vendors needn't publish thick, glossy brochures to hand out to visitors. Instead, they may print various stylish promo cards for distribution. When later presented to a Bedoop sensor, each card leads to a web-based presentation – optionally including persuasive video and other multi-media components. The user can be prompted to provide data to customize, or focus, the presentation to the user's particular requirements. If the user wants further information, a request can be made by the click of a mouse (or the twist of a card).

Prizes and Product Promotions

Product packaging (e.g., Coke cans, Snapple bottles, Pepsi 12-pack boxes) can be encoded for contest purposes. The encoding can be customized, item to item, so that selected items – when Bedoop scanned – are recognized to be the one in a hundred that entitles the owner to a cash or merchandise prize. A remote server to which the item's Bedoop data is provided queries the user for contact information (e.g., address, phone number) so the prize can be awarded or, for smaller prizes, the system can print out an award certificate redeemable at local merchants for products or cash. Once a winning item is identified to the remote server, its UID on the server is marked as redeemed so that the item cannot later be presented to win another prize.

In other such embodiments, all of the items are encoded identically. Winners are determined randomly. For example, during a contest period, persons around the world may present Coke cans to Bedoop systems. The corresponding Bedoop application on each user computer submits Bedoop data to a corresponding web address. The user's e-mail address

may also be included with the submission. As this data is relayed to the corresponding server computer(s), every Nth set of data is deemed to be a winner, and a corresponding award notification or prize is dispatched to the Bedoop system from which the winning set of data originated.

5 The server computer that receives such contest submittals from client Bedoop systems can be arranged to prevent a single user from bombarding the server with multiple sets of data in an attempt to win by brute force. (This may be done, for example, by checking the included e-mail address, and not considering a data submittal if the same e-mail address was encountered in data submitted within the past hour. Similar anti-brute-force protection can be
10 provided on the user's computer, preventing, e.g., repeated contest data to be sent more frequently than once per hour. More sophisticated anti-brute-force measures can of course be provided.)

Product Information and Ordering

15 In accordance with another embodiment of the present invention, product packaging and product advertisements can be encoded with Bedoop data that, when presented to a Bedoop system, initiates a link to a web page from which that product can be purchased, or more information obtained. Once the link has been established, the user can be instructed to manipulate the object in different of the earlier-described modes to effect different functions,
20 e.g., move towards camera to order the product; move away from camera for product information. If the object is moved towards the camera to effect an order, the user can be prompted to further manipulate the object to specify delivery options (e.g., rotate left for overnight mail, rotate right for regular mail). If the object is moved away from the camera to request product information, the user can be promoted to further manipulate the object to
25 specify the type of information desired (e.g., rotate left for recipes, rotate right for FDA nutritional information, move up for information on other products in this family, move down to send an email to the product manufacturer).

 Credit card or other customer billing information, together with mailing address information, can be stored in a profile on the Bedoop system, and relayed to the transactional
30 web site either automatically when a purchase action is invoked, or after the user affirms that such information should be sent (which affirmation may be signaled by manipulation of the

packaging or advertisement in one of the earlier-described modes). Other modes of payment can naturally be employed. (One such alternative is the first-to-redeem electronic money system described in the present assignee's patent application 60/134,782.)

5 Clothing

In accordance with another aspect of the invention, clothing can be ordered on-line by presenting to a Bedoop system a photograph from a catalog, or a garment tag or label. Encoded on each is product-identifying data, including a manufacturer ID. The Bedoop system responds by establishing a link to a remote computer maintained by or on behalf of the manufacturer. In addition to relaying the product identification data to the remote computer, the Bedoop application also sends some or all of a clothing profile maintained by the user on the local computer. This profile can specify, e.g., the person's weight, height, shoe size, waist size, inseam, etc. The remote computer can confirm availability of the identified item in the size specified in the clothing profile, and solicit payment and shipping instructions.

15 Computer Access Cards

This disclosure earlier considered access cards used to gain access to secure buildings. Related principles can be used in conjunction with computer access.

A driver's license, employee photo ID, or other such document can be presented to a Bedoop sensor on a computer. The computer recognizes the user and can take various steps in response.

One response is to log onto a network. Another is to set load a user profile file by which the computer knows how to arrange the desktop in the user's preferred manner. By manipulating the Bedoop-encoded object, the user can further vary the environment (e.g., rotate left to launch standard business productivity applications and software development applications; rotate left to launch lunchtime diversions – stock update, recreational games, etc.)

Hotel rooms are increasingly providing computer services. By presenting a driver's license, a Bedoop-equipped computer in a hotel room can link to a remote site indicated by the Bedoop data, obtain preference data for that user, and launch applications on the hotel computer in an arrangement that mimics that user's familiar work computer environment.

Audio/Video Disks, Software, and Books

Bedoop data can be conveyed by indicia or texturing on the surfaces of CD and DVD disks, on the labels (or authenticity certificates) for same, on the enclosures for same (e.g., jewel box, plastic case, etc.), on book dust jackets, on book pages, etc. Any of these objects can be presented to a Bedoop device to establish a link to a related web site. The consumer can then manipulate the object (or otherwise choose) to select different options.

For music, one option is to receive MP3 or other clips of songs by the same artist on other CDs, or of songs from other artists of the same genre. Another is to view music video clips featuring the same artist. Still another is to order tickets to upcoming concerts by that artist. In-store kiosks can permit tentative customers to listen to sample tracks before they buy.

Similar options can be presented for video DVDs. In the case of video, this can include listings of other movies with the same director, with the same star(s), etc. In the case of software, the options can include advisories, bug fixes, product updates and upgrades, etc. Naturally, the user can make purchases from these sites, e.g., of other music by the same artist, other videos with the same star, software upgrades, etc.

Similar options can be accessed using Bedoop data associated with printed book materials.

Ad Tracking

Advertisers commonly use different advertisements for the same product or service, and employ means to track which ad is more effective within which demographic group. Bedoop can provide such functionality.

Consider a travel service web site that is promoting Hawaiian vacations. Bedoop data from several advertisements can lead consumers to the site.

Identical advertisements can be placed in several different magazines. Each is encoded with a different Bedoop UID. By monitoring the UIDs of the Bedoop inquiries to the site, the travel service can determine which magazines yield the highest consumer response (e.g., per thousand readers).

Likewise, within a single magazine, two or more advertisements may be encoded with Bedoop data leading to the site – again, each with a different UID. Again, analysis of the UIDs used in accessing the site can indicate which advertisement was the more effective.

The instantaneous nature of the internet links permits advertisers to learn how consumer responses to print advertisements vary with time-of-day, yielding information that may assist in making ads for certain products more effective.

More elaborate variants and combinations of the foregoing are, of course, possible. If the consumers provide personal information in response to the ads (either by permitting access to pre-stored personal profile data, or by filling in web-based forms, or by manipulation of the ad (e.g., “please move the ad towards your Bedoop sensor if you drank coffee this morning”)), still richer statistical data can be gleaned.

Rolodex of Cards

Bedoop-encoded business cards as detailed above can be accumulated and kept near a telephone or computer in a Rolodex-like arrangement. If a refrigerator ice-maker malfunctions, a homeowner can find the card for the appliance repairman used a few years ago, and present it to a Bedoop sensor. A link is established to the repairman’s company (e.g., web site or via telephone). At a web site, the repairman may provide basic information, such as hours of availability, current fee schedule, etc. The homeowner may select an option (by card gesture or otherwise) to invoke a teleconference (e.g., NetMeeting) to consult about the problem. Or the homeowner may select another option to send e-mail. Still a further option may permit the homeowner to schedule a house call on the repairman’s weekly calendar. Still a further option may permit the homeowner to view one or more short videos instructing customers how to fix certain common appliance problems.

Stored Value Cards

The earlier cited “first-to-redeem” electronic money system may encode Bedoop data on a card that leads to storage at which the random-number tokens (which represent increments of money) are stored. Presenting the card to a Bedoop system launches an application that reads and encrypts the tokens and forwards the encrypted data to the clearinghouse computer of the corresponding bank to learn their remaining value. There the

tokens are decrypted and checked for validity (but not redeemed). The bank computer responds to the Bedoop system, indicating the remaining value of the tokens on the card.

For security reasons, the storage containing the random-number tokens should not be generally accessible. Instead, the user must provide authentication data indicating authorization to gain access to that information. This authentication data may be a PIN code. Or the user may provide authentication by presenting a second Bedoop-encoded object, e.g., a driver's license to the Bedoop system. (Many other Bedoop systems may advantageously use, or require the use of, two or more Bedoop objects – either presented one after the other, or all at the same time. The Bedoop system can provide visual or audible prompts leading the user to present the further Bedoop object(s) as necessary.

Ski Lift Tickets

In accordance with another embodiment, ski lift tickets are Bedoop encoded to provide various functionality.

For example, instead of buying a lift ticket good for a day, a skier may purchase a ticket good for eight lifts. This data is encoded on the ticket, and sensed by a Bedoop sensor at each lift. The sensors are networked to a common server that tracks the number of lifts actually purchased, and updates the number as used. The skier is informed of the number of rides remaining on entering or leaving the lift. Statistical data can be collected about trail usage (e.g., N% percent of skiers ski all day along just two lifts, etc.).

Off the slopes, back at home, the used lift ticket may be presented to a Bedoop sensor to obtain current snow conditions and lift hours, or to review trail maps, or to order ski vacation packages. If the ticket is encoded with the owner's name, UID, or other information of commercial/marketing interest, local merchants may give the bearer discounts on selected goods in response to Bedoop scanning of the ticket and recovery of such information.

REI Membership Cards

Membership cards for certain stores can be Bedoop-encoded to provide added value to the member. For outdoor gear stores such as REI, presentation of the card to a Bedoop sensor can lead to a library of USGS maps, to web pages with current fishing and hunting regulations, etc. Naturally, the store's on-line ordering site is just a quick twist away.

Theme Park Tickets

Theme park tickets can be encoded with the age and gender of the visitor, and with additional data permitting the experience to be customized (e.g., from a roster of theme park personalities, the visitor's favorite is Indiana Jones). Throughout the park are kiosks to which the visitor can present the ticket to orchestrate the visit to follow a particular story line. Some kiosks issue premiums matching the age/gender of the recipient.

Car Keys

In accordance with another embodiment of the invention, car keys (or key ring fobs) are Bedoop encoded. When the car is taken to a shop for service, the mechanic presents the key to a Bedoop sensor, and thereby obtains the car's maintenance history from a remote server on which it is maintained. At home, the key can be presented to a Bedoop sensor and manipulated to navigate through a variety of automotive-related web sites.

In some embodiments, the Bedoop-encoded object is not used to navigate to a site, but is instead used to provide data once a user's computer is otherwise linked to a web site. A user surfing the web who ends up at a car valuation site can present a key to the Bedoop scanner. The Bedoop data is used to access a remote database where the make, model, options, etc., of the car are stored. This data is provided to a database engine that returns to the user the estimated value of the car.

While visiting a mechanic's web site, presentation (and optionally manipulation) of a key or key ring fob can be employed to schedule a service appointment for the car.

Fashion Coordination

Some department stores and clothing retailers offer "personal shoppers" to perform various services. For example, a customer who is purchasing a dress may ask a personal shopper for assistance in selecting shoes or accessories that complement the dress.

A Bedoop-encoded garment tag on the dress can be employed to obtain similar assistance. In response to such a tag, a Bedoop system can query a database to obtain a mini-catalog of clothes and accessories that have previously been identified as complementing the dress identified by the tag. These items can be individually displayed on a screen associated

with the system, or a virtual model wearing the dress - together with one or more of the recommended accessories - can be synthesized and depicted. The shopper may quickly review the look achieved by the model wearing the dress with various different pairs of shoes, etc., by repeatedly activating a user interface control (by mouse, touch screen, or garment tag gestures) to cycle through different combinations.

A shopper's credit card can be Bedoop-encoded so as to lead Bedoop systems of particular stores (i.e., stores pre-authorized by the shopper) to a profile on the shopper (e.g., containing size information, repeat purchase information, return history, style/color preferences, etc.).

Credit Card Purchases

When a consumer visits a commercial web site and wishes to purchase a displayed product, the transaction can be speeded simply by presenting a Bedoop-encoded credit card to a Bedoop sensor on the user's computer. The Bedoop data on the card leads to a database entry containing the credit card number and expiration date. The Bedoop application then sends this information (optionally after encrypting same) to the web site with instructions to purchase the depicted product.

(Impulse purchases are commonly deterred by the hurdles posed between the purchase impulse and the completed purchase. This and other Bedoop applications aid in reducing such hurdles.)

Product Marketing

Bedoop data relating to one product or service can be used to cross-market others products and services. Consider a consumer who purchases a pair of golf shoes. The box is Bedoop encoded. By presenting the box to a Bedoop system, the consumer is linked to a web page that presents various promotional offers. The consumer may, for example, elect to play a free round of golf at one or more identified local golf courses, or print a coupon for ten percent off any order of socks from an on-line sock merchant. (Various means can be employed to prevent multiple redemptions from a single box. One is a serial number that is tracked by the web page or cross-marketed merchant, and only honored once. Another is

identification data corresponding to the consumer that is tracked to prevent multiple redemptions.)

Product tags can likewise be Bedoop-encoded. A tag from an article of Nike apparel can lead to the Nike on-line store, where the user can buy more merchandise. If the tag is from a soccer jersey, a certain tag manipulation (e.g., rotate left) may lead the user to a special-interest soccer page, such as for the World Cup. A tag on a golf glove may lead to a website of a local golf course. Twist left to reserve a tee time; twist right to review course maps and statistics. Bedoop kiosks can be provided in retail stores to let consumers use the Bedoop features.

Travel Planning Services

After making a reservation at a resort, a consumer is typically mailed (by email or conventional mail) various confirmation information. If not already printed, the consumer can print this information (e.g., a confirmation card).

Bedoop-encoding on the printed object can lead to web-based information relating to the reservation (e.g., reservation number, the consumer's name, arrival/departure dates, etc.). If the consumer wishes to make dinner or golf reservations, this object is presented to a Bedoop system – either at the user's home, at an airport kiosk, etc. The system recognizes the object type and encoded data, and establishes a link to a remote computer that provides various information and scheduling services for the resort. By manipulating the object (or otherwise) the consumer selects desired dinner and golf tee times. The system already has the reservation number (indexed by the UID), so tedious provision of such data is avoided.

In some embodiments, the remote computer is not maintained by the resort, but is rather maintained by an independent travel service. (The travel service may also maintain the DNS leaf node server.) The computer can present a web page (branded by the travel service or not) that offers the scheduling options desired by the user, and also presents links to other information and services (e.g., offering entry tickets to nearby attractions, and advertising nearby restaurants).

Airline tickets (or e-ticket confirmations) can be similarly encoded with Bedoop data. These items may be presented to Bedoop systems – at a traveler's home or in airports – to permit review and changing of travel itinerary, reserve hotels and rental cars, secure first-class

upgrades, check the airplane's seating arrangement, review frequent flier status, scan tourist information for the destination, etc.

Movie Tickets

5 As indicated earlier, movie tickets can be encoded with Bedoop data identifying, e.g., the movie title and date. When a movie viewer returns home, the ticket stub can be presented to a Bedoop system. One of the options presented by the corresponding Bedoop application can be to launch a pay-per-view screening of the just-seen movie at a discounted rate. Another is to download the movie onto a writeable DVD disk at the viewer's home, perhaps serialized
10 to permit playback only on that viewer's DVD player, or enabled for only a few playbacks, etc. (again, likely for a discounted fee). Still another option is to present web-delivered video clips from the movie. Another is to offer related merchandise for purchase, possibly at discount to retail. (These features may be available for only a limited period after the date encoded on the ticket stub.) Another is to alert the consumer to upcoming movies of the
15 same genres, or with the same director or stars, or released by the same studio. Still another is to direct a web browser to an on-line ticket merchant for tickets to other movies. The consumer may navigate among these options by manipulating the ticket stub, or otherwise.

 The same, or related, options can likewise be provided in response to Bedoop data detected from a book jacket presented to a Bedoop system.

Video Recording

 A video recording device can be programmed to record a broadcast program by presenting a Bedoop sensor with a printed promotion for the program (e.g., an advertisement in a newspaper or TV Guide). Bedoop-encoded within the printed document is data by which
25 the Bedoop system (which may be built into the video recorder or separate) can set the recording time, date, and channel.

Set Top Boxes

 Many entertainment-related applications of Bedoop data can be implemented using
30 television set top boxes. Such boxes include processors, and typically include a return

channel to a control facility. The provision of a Bedoop chip and optical sensor can vastly increase the functionality these devices presently provide.

Special Event Tickets

5 Consider a ticket to a basketball game. By presenting the ticket to a Bedoop system, a user may access the web site of either team so as to review recent scores and statistics. The user may also obtain a web-based virtual tour of the arena, and review seating maps. Tickets for upcoming games may be ordered, as well as pay-per-view games and team souvenirs. For high-priced tickets, the user may be entitled to premium web features, such as on-line text-,
10 audio-, or video-chat session with a team star on the day before the game.

 Unlike conventional tickets, Bedoop-encoded tickets need not limit the user to a predetermined seat. While the ticket may be printed with a nominal seat, the user may present the ticket to a Bedoop sensor and access a web site at which a different seat can be reserved. On attending the event, the consumer presents the ticket to a Bedoop sensor that reads the
15 ticket UID and looks up the seat assignment most-recently picked by the consumer. It then prints a chit entitling the consumer to take the seat earlier selected from the transactional web site.

Signet Rings

20 Signet rings have historically been used to indicate a person's identity or office. Such rings, or other items of personal jewelry, can be encoded with Bedoop data (either by texturing or printing) and presented as necessary to Bedoop systems. The extracted Bedoop data can lead to a secure web site indicating the person's name and other information (i.e., a web site that has anti-hacking measures to prevent illicit change of the stored identification
25 information). Such a signet ring can be presented to Bedoop systems that require a high-confidence confirmation of identity/authorization before proceeding with a Bedoop function.

Post-It[®] Notes

 Pads of Post-It[®] notes, or other pads of paper, can be marked by the manufacturer
30 (either by texturing, watermarked tinting, ink-jet spattering, etc.) to convey steganographic

data (e.g., Bedoop data). When such a note is presented to a Bedoop system, the system may launch an application that stores a snapshot of the note. More particularly, the application may mask the note-portion of the image data from the other image data, virtually re-map it to a rectangular format of standardized pixel dimensions, JPEG-compress the resulting image, and store it in a particular computer subdirectory with a name indicating the date of image acquisition, together with the color and/or size of the note. (These latter two data may be indicated by data included in the Bedoop payload.) If the color of the note is indicated by digital data (e.g., in the file name), then the image itself may be stored in grey-scale. When later recalled for display, the white image background can be flooded with color in accordance with the digital color data.

The Bedoop system may buffer several past frames of image data. When the object is recognized as a Post-It note whose image is to be saved, the system may analyze several such frames to identify the one best-suited for storage (e.g., check the spatial frequency content of the note as imaged in each frame, to identify the one with the finest detail), and store that one.

When a Post-It note is recognized by the Bedoop system, the system may emit a confirmation tone (or other response) to indicate that the object has been recognized, but not immediately execute the snapshot operation. Instead, the system may await a further instruction (e.g., gesture) to indicate what operation is desired.

By moving the note towards the sensor, for example, the user can signal that a snapshot operation is to be performed. (This closer presentation of the note may also permit the imaging system to capture a more detailed frame of image data.)

By moving the note away, the system may respond by reading, decompressing, and displaying the six most-recently stored Post-It note images, in tiled fashion, on the computer screen. The individual notes can be displayed at their original dimensions, or each can be resized to fill the full height or width of a tile. A user interface control (responsive to gestures, mouse operation, keyboard scroll arrows, etc.) allows the user to scroll back in time to any desired date.

The full 64-bit Bedoop payload of other embodiments may not be needed for Post-It notes. In the just-given example, for example, the Bedoop system responds to all Post-It notes in the same fashion. Thus, an abbreviated Bedoop format that indicates simply 'I'm a Post-It note, yellow, size 3" x 3"' can suffice. The twelve bit CLASS ID, with eight further

bits to indicate color/size combinations, may be sufficient. Reducing the payload permits it to be more robustly encoded on small objects. (As noted below, Bedoop decoding systems can look for several different data formats/protocols in trying to extract Bedoop data from an object.)

5

Alignment of Documents for Other Purposes

While the just-described pre-marked paper triggered a Bedoop response when presented to a Bedoop sensor (i.e., take a snapshot of the paper), the markings can be used for purposes other than to trigger Bedoop responses.

10

Regardless of the particular data with which the paper is encoded, the embedded subliminal graticules, or other steganographically-encoded registration data, can be used by other applications to correct misalignment of scanned data. In a photocopier, for example, a document need not be placed exactly squarely on the glass platen in order to yield a properly-aligned photocopy. The scanner scans the skewed document and then detects the

15

steganographic registration markings in the resulting scan data. This data is then processed to virtually re-register same, so that the registration markings are in a desired alignment. The processed scan data is then provided to the xerographic reproduction unit to yield a photocopy in which the skew effect is removed.

20

The same technique is likewise applicable to video recorders, digital cameras, etc. If such a device images an object (e.g., a photograph) with steganographic registration markings, these markings can be used as a guide in re-registering the resulting data to remove mis-alignment effects.

Postal Mail Information

25

Many contexts arise in which data to be presented to a consumer is valuable only if timely. The postal service mail is ill-suited for some such information due to the latency between printing a document, and its ultimate delivery to a recipient. Bedoop principles, however, allow the recipient to take a postal object that was printed well before delivery, and use it on receipt (i.e., present to a Bedoop system) to receive up-to-the-minute information. In this and other embodiments, the Bedoop data can also uniquely identify the

30

Distributors of printed advertising can reward Bedoop-driven consumer visits to their web sites by issuing digital tokens or coupons that can be redeemed for premiums, cash-back, etc. Every millionth visitor wins a million pennies (with appropriate safeguards, e.g., preventing more than one entry an hour).

5

Classes of Bedoop Encoding

The above-described embodiments focused on use of Bedoop data after decoding. Additional insight may be gained by examining the earlier part of the process – encoding.

Encoding can be performed in many contexts, which may be conceptualized as falling into three broad classes. The first is static marking, in which a document designer, pre-press service bureau, advertising agency or the like embeds Bedoop data. The second is dynamic marking, in which automated systems encode, or vary, Bedoop data “on the fly.” Such systems can tailor the Bedoop data to particularly suit the context, e.g., to the moment, place, user, etc. The third is consumer marking, in which Bedoop data is added to a document at the time of printing.

15

The second class of encoding enables features not available from the first. Consider an American Express travel web page with information about travel to Hawaii. A DNS leaf node server points to this page in response to certain Bedoop data – e.g., data encoded in a magazine photograph of a Hawaiian beach scene.

20

Actually, all Bedoop data having a certain CLASS and DNS ID may lead to this web page, irrespective of the UID data. If the magazine photo is encoded with a particular “don’t care” UID field (e.g., 11111111111111111111111111111111), this may signal the originating Bedoop system – or any intervening system through which the Bedoop data passes – that arbitrary data can be inserted in the UID field of that Bedoop packet. The originating Bedoop system, for example, can insert a dynamically-configured series of bits into this field. Some of these bits can provide a profile of the user to the remote server, so that the Bedoop response can be customized to the user. (The user would naturally pre-approve information for such use so as to allay privacy concerns.)

25

As one example, the local Bedoop system can set the least significant bit of the UID field to a “0” if the user is male, or to a “1” if the user is female. The next four bits can

30

indicate the user's age by one of sixteen age ranges (e.g., 3 or less, 4-5, 6-7, 8-9, 10-11, 12-13, 14-15, 16-17, 18-20, 21-24, etc.).

Alternatively, or in addition, the local Bedoop system can stuff the don't-care UID field (all of it, or in part) with signature data tending to uniquely identify the local Bedoop system (e.g., system serial number, a hash code based on unchanging data unique to that system, etc.) By reference to such data, the remote server can identify repeat visits by the same user, and can tailor its responses accordingly (e.g., by recalling a profile of information earlier entered by the user and stored at the remote server, avoiding the need for data re-entry).

More on Optical Input Devices

It is expected that image input devices will soon become commonplace. The provision of digital cameras as built-in components of certain computers (e.g., the Sony Vaio laptops) is just one manifestation of this trend. Another is camera-on-a-chip systems, as typified by U.S. Patent 5,841,126 and detailed in Nixon et al., "256x256 CMOS Active Pixel Sensor Camera-on-a-Chip," IEEE J. Solid-State Circuits, Vol. 31(12), pp. 2046-2051 (1996), and Fossum, "CMOS Image Sensors: Electronic Camera-on-a-Chip," IEEE Transactions of Electron Devices, vol. 44, No. 10, Oct. 1997. Still another is head-mounted cameras (as are presently used in some computer-augmented vision systems). These and other image input devices are all suitable for use in Bedoop systems.

Camera-on-a-chip systems can be equipped with Bedoop detector hardware integrated on the same chip substrate. This hardware can be arranged to find and decode Bedoop data from the image data – notwithstanding scale, rotation, differential scaling, etc. Gestural decoding can also be provided in hardware, with the resulting data output in packet form on a serial output bus. Such a chip can thus provide several outputs – image data (either in raw pixel form, or in a data stream representing the image in one of various image formats), 64 bits of Bedoop data (serially or in parallel), and decoded gesture data.

In other embodiments, the Bedoop detector (and/or the gestural decoder) can be on a substrate separate from the camera system.

To accommodate different Bedoop data formats and protocols, the hardware can include RAM or ROM in which different format/protocol information is stored. (These

different formats/protocols can relate, e.g., to Bedoop systems employing different data payload lengths, different subliminal grids, different encoding techniques, etc.) As the Bedoop system stares out and grabs/analyzes frames, each frame can be analyzed in accordance with several different formats/protocols to try and find a format/protocol that yields valid Bedoop output data.

Movable Bedoop Sensors

Although the illustrated Bedoop systems are generally stationary, they need not be so. They can be portable. Some such systems, for example, employ palmtop computers equipped with optical sensor arrays. If the palmtop is provided with live network connectivity (e.g., by wireless), then Bedoop applications that rely on remote computers can be implemented just as described. If the palmtop is not equipped with live network connectivity, any Bedoop applications that rely on remote computers can simply queue such communications, and dispatch same when the palmtop next has remote access (e.g., when the palmtop is next placed in its recharger and is coupled to a modem through which internet access can be established).

Another variant is a Bedoop sensor that is movable around a desk or other work-surface, like a mouse. Such a sensor can be coupled to the associated computer by cabling, or a wireless interface can be used. The peripheral may be arranged for placement on top of an item in order to read digital data with which the object is marked. (Built-in illumination may be needed, since the device would likely shadow the encoding.) Some forms of such peripherals are adapted to serve both as general purpose digital cameras, and also as Bedoop sensors.

Such a peripheral would find many applications. In “reading” a magazine or book, for example, it may be more intuitive to place a Bedoop reader “on” the object being read, rather than holding the object in the air, in front of a Bedoop sensor. This is particularly useful, e.g., when a magazine page or the like may have several differently-encoded Bedoop sections (corresponding to different articles, advertisements, etc.), and the user wants to assure that the desired Bedoop-encoded section is read.

The “bookmark” paradigm of internet browsers might be supplemented with paper bookmarks, e.g., Bedoop data encoded on one or more pages of paper. To direct a browser to

a particular bookmarked destination, the peripheral is simply placed on top of the page (or part thereof) that is marked with the corresponding Bedoop data. A user may print a "map" comprised of postage stamp-sized regions tiled together, each of which regions represents a favorite web destination.

5 Such a map may be printed on a mouse pad. Indeed, mouse pads with certain maps pre-encoded thereon may be suitable as promotional materials. A company may offer to print a family photograph on such a pad. Encoded within the photograph or the pad texture are addresses of web sites that have paid a fee to be accessible in this manner on a user's desk.

10 Like mice -- which are provided with buttons, roller wheels, and roller buttons in addition to X-Y encoders -- movable Bedoop encoders can likewise be provided with auxiliary switches and roller inputs to complement the data input provided by the optical sensor. Indeed, some embodiments integrate the functions of Bedoop peripheral with a mouse. (The undersides of mice are generally under-utilized, and can readily be equipped with an image sensor.)

15 Gestural input can readily be provided by such a peripheral -- in this context moving the sensor rather than the object.

Watermarking Techniques

20 There are nearly as many techniques for digital watermarking (steganographic data encoding) as there are applications for it. The reader is presumed to be familiar with the great variety of methods. A few are reviewed below.

mu 057
25 The present assignee's prior application 09/127,502, filed July 31, 1998, shows techniques by which very fine lines can be printed on a medium to slightly change the medium's apparent tint, while also conveying digital data. Commonly-owned application 09/074,034, filed May 6, 1998, details how the contours of printed imagery can be adjusted to convey digital data. (That technique can be applied to printed text characters, as well as the line art imagery particularly considered.) The assignee's patent 5,850,481 details how the surface of paper or other media can be textured to convey optically-detectable binary data. The assignee's patents 5,841,886 and 5,809,160 detail various techniques for
30 steganographically encoding photographs and other imagery.

Some watermarking techniques are based on changes made in the spatial domain; others are based on changes made in transformed domains (e.g., DCT, wavelet).

Watermarking of printed text can be achieved by slight variations to character shape, character kerning, line spacing, etc.

5 Data glyph technology, as detailed in various patents to Xerox, is usable in many of the applications detailed herein.

The foregoing is just a gross under-sampling of the large number of watermarking techniques. The artisan is presumed to be familiar with such art, all of which is generally suitable for use in the applications detailed herein.

10 More generally, essentially any data encoding method that permits recovery of the encoded data from optical scan data can be employed. Bar codes (1D and 2D) are but the most familiar of many such optically-detectable data encoding techniques.

Conclusion

15 Having described and illustrated the principles of our invention with reference to illustrative embodiments, it should be recognized that the invention is not so limited.

For example, while certain of the embodiments were illustrated with reference to internet-based systems, the same techniques are similarly applicable to any other computer-based system. These include non-internet based services such as America Online and
20 Compuserve, dial-up bulletin board systems, etc. Likewise, for internet-based embodiments, the use of web browsers and web pages is not essential; other digital navigation devices and other on-line data repositories can be similarly accessed.

Similarly, while the details of the preferred Bedoop system were particularly given, the underlying principles can be employed in numerous other forms.

25 For example, one other form is to steganographically encode physical objects with Digital Object Identifiers (DOIs). The Center for National Research Initiatives and the Digital Object Identifier Foundation (www.doi.org) have performed extensive work in establishing an infrastructure by which digital objects can be distributed, tracked, and managed. Some of this same infrastructure and technology can be adapted, in accordance with
30 the teachings provided above, to associate new functionality with physical objects.

Another form is not to reference a remote data repository by data embedded on an object, but instead to encode the ultimate data directly on the object. A photograph, for example, can be literally encoded with a telephone number. On presenting the photograph to an optical sensor on the telephone, the telephone can analyze the optical information to extract the telephone number, and dial the number, without the need for any external data. Similarly, a printed office document (e.g., spreadsheet) can be encoded with the path and file name of the corresponding electronic file, obviating the need for indirect linking (e.g., to a database to correlate a UID to a computer address). Most of the above-described embodiments are suitable for such direct encoding of the related data.

In the business card example given above, the detailed techniques can be supplementary to existing optical character recognition techniques. That is, the image data from an optical sensor can be applied both to a Bedoop decoder and to an OCR system. Text characters discerned by the OCR system can be entered directly into a contacts manager personal database. The techniques employed in the Bedoop system to locate the encoded object and handle visual distortion (e.g., the visual artifacts due to scale, rotation, etc.) can advantageously be used in OCR detection as well, permitting extraction of the OCR information without careful placement of the card.

While certain of the foregoing embodiments made reference to ink-jet printing, similar advantages can often be obtained with other printing technologies, e.g., laser/xerographic printing, offset printing, etc.

In the foregoing embodiments, Bedoop decoding generally proceeded from image data obtained from a physical object. However, in some contexts, it is advantageous to Bedoop-decode image data provided electronically, e.g., over the internet.

Likewise, while the foregoing embodiments generally relied on Bedoop image sensors that stared out for an object at an expected point, in alternative embodiments, sensors that seek rather than stare can be employed (as was illustrated above in connection with the elevator example)

Similarly, while the illustrated embodiments generally employed sensors that repeatedly grabbed frames of image data, this need not be the case. Single frame systems, such as flatbed scanners, and video systems arranged to grab single frames -- with or without TWAIN interfaces -- can alternatively be used.

As indicated above, while steganographic encoding of the digital data is used in the preferred embodiments, visible forms of digital encoding – such as bar codes – can naturally be employed where aesthetic considerations permit.

In certain of the embodiments, digital data conveyed by means other than optical can be used. Electromagnetic detection (e.g., of the sort used in proximity-based card-access systems) can be arranged to decode digital data, permitting “at-a-distance” reading of data from physical objects, just as in the foregoing embodiments.

Since the Bedoop image sensors typically acquire plural frames of data, the extraction of the digital data can be based on more than a single image frame. More confidence in the results may be accumulating decoded data over several frames. Moreover, movement of the object within the sensor’s field of view may permit the system to acquire information from other perspectives, etc., enhancing system operation.

While the preferred embodiments employ 2-D image sensors (e.g., CCDs), other optical sensing technology can alternatively be employed. Supermarket laser scanners, for example, can read bar-code data. Raster-scanning of such systems can permit acquisition of 2-D data (either in bit-mapped form, or grey-scale).

While the illustrated embodiments used a 12/24/24 bit protocol for CLASS/DNS/UID data, other arrangements can of course be used. In some applications it is advantageous for the protocol to more nearly match those commonly used for internet communications. For example, IP addresses for internet Domain Name Servers (DNS) are presently 32 bits, with extension to 64 or 128 bits foreseen in the near future. The DNS field in Bedoop systems can be follow the internet standard.

Some embodiments can advantageously employ texture-based Bedoop encoding of objects. Bedoop texturing can be effected by various means, including pressure rollers, chemical or laser etching, etc.

While the foregoing embodiments have generally employed planar objects to convey the digital encoding, this need not be the case. Objects of other shapes can likewise be employed. Some shapes present relatively straightforward image processing tasks. Data imaged from a soft drink can or other cylindrical surface, for example, is fairly easy to remap using known geometrical transforms so as to essentially “unwrap” the printing from the can. Other geometries can present more complex re-mappings, but are likewise generally within

the capabilities of the artisan. (Such remapping is facilitated by encoding in the data certain reference markings, such as subliminal graticules, etc. The unknown 3D shape of the object being imaged can usually be inferred from the apparent warping of the reference markings in the 2D image data generated by the scanner. Once the warping is characterized, it is generally straightforward to un-warp so as to prepare the image data for decoding.)

It was once popular to predict that paper documents would be replaced with electronic media. In hindsight, electronic media may be recognized as a poor surrogate for paper. Electronic media conveys information flawlessly, but is lacking in experiential attributes. We can hold paper, stack it, own it, deface it, give it, guard it, etc. It provides an opportunity for physical dominion entirely lacking with electronic media.

From the foregoing discussion it can be seen that, rather than replacing paper with electronic media, perhaps the future lies in giving paper digital attributes – hybridizing the physical experience of paper with the technical advantages of digital media. Such an arrangement makes available a great wealth of new functionality, now accessible through familiar paper items, rather than through a “computer input peripheral.”

To provide a comprehensive disclosure without unduly lengthening this specification, applicants incorporate by reference the patents, applications, and publications identified above.

In view of the many embodiments to which the principles of our invention may be applied, it should be recognized that the detailed embodiments are illustrative only and should not be taken as limiting the scope of our invention. Rather, we claim as our invention all such embodiments as fall within the scope and spirit of the following claims, and equivalents thereto.